AFRL-RI-RS-TR-2012-135

CLUSTERING THEORY AND APPLICATIONS

CREATIVE STEP LLC

APRIL 2012

FINALTECHNICAL REPORT

APPROVED FOR PUBLIC RELEASE; DISTRIBUTION UNLIMITED.

STINFO COPY

AIR FORCE RESEARCH LABORATORY
INFORMATION DIRECTORATE

■ **AIR FORCE MATERIEL COMMAND** ■**UNITED STATES AIR FORCE** ■ **ROME, NY 13441**

NOTICE AND SIGNATURE PAGE

Using Government drawings, specifications, or other data included in this document for any purpose other than Government procurement does not in any way obligate the U.S. Government. The fact that the Government formulated or supplied the drawings, specifications, or other data does not license the holder or any other person or corporation; or convey any rights or permission to manufacture, use, or sell any patented invention that may relate to them.

This report is the result of contracted fundamental research deemed exempt from public affairs security and policy review in accordance with SAF/AQR memorandum dated 10 Dec 08 and AFRL/CA policy clarification memorandum dated 16 Jan 09. This report is available to the general public, including foreign nationals. Copies may be obtained from the Defense Technical Information Center (DTIC) (http://www.dtic.mil).

AFRL-RI-RS-TR-2012-135 HAS BEEN REVIEWED AND IS APPROVED FOR PUBLICATION IN ACCORDANCE WITH ASSIGNED DISTRIBUTION STATEMENT.

FOR THE DIRECTOR:

/ s /	/ s /
MICHAEL O. LITTLE	PAUL ANTONIK, Technical Advisor
Work Unit Manager	Computing & Communications Division
	Information Directorate

This report is published in the interest of scientific and technical information exchange, and its publication does not constitute the Government's approval or disapproval of its ideas or findings.

REPORT DOCUMENTATION PAGE

Form Approved
OMB No. 0704-0188

1. REPORT DATE *(DD-MM-YYYY)*	2. REPORT TYPE	3. DATES COVERED *(From - To)*
APR 2012	Final Technical Report	JUN 2010 – OCT 2011

4. TITLE AND SUBTITLE

CLUSTERING THEORY AND APPLICATIONS

5a. CONTRACT NUMBER FA8750-10-2-0115

5b. GRANT NUMBER N/A

5c. PROGRAM ELEMENT NUMBER 61102F

6. AUTHOR(S)

H. T. Kung

5d. PROJECT NUMBER 23TA

5e. TASK NUMBER MH

5f. WORK UNIT NUMBER 01

7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES)
Creative Step LLC
24 Summit Road
Belmont, MA 02478

8. PERFORMING ORGANIZATION REPORT NUMBER
N/A

9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)

Air Force Research Laboratory/RITB
525 Brooks Road
Rome NY 13441-4505

10. SPONSOR/MONITOR'S ACRONYM(S)
AFRL/RI

11. SPONSORING/MONITORING AGENCY REPORT NUMBER
AFRL-RI-RS-TR-2012-135

12. DISTRIBUTION AVAILABILITY STATEMENT
Approved for Public Release; Distribution Unlimited. This report is the result of contracted fundamental research deemed exempt from public affairs security and policy review in accordance with SAF/AQR memorandum dated 10 Dec 08 and AFRL/CA policy clarification memorandum dated 16 Jan 09.

13. SUPPLEMENTARY NOTES

14. ABSTRACT
The report describes the theory for recovering information from sparse signal representations in distributed sensing applications. This theory is useful in streamlining networking and decoding operations over bandwidth constrained wireless networks. Using these results, we can for example transport and store a single combined measurement set, rather than multiple sets from all sensors. We show that via source separation and joint decoding, it is possible to recover information about the original signal from such combined measurements using progressive reconstruction. This results indicate a substantial reduction in the number of variables that are decoded at each step with a much reduced decoding time.

15. SUBJECT TERMS

Compressive Sensing, Signal Recovery, Clustering, Target Detection

16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT	18. NUMBER OF PAGES	19a. NAME OF RESPONSIBLE PERSON
					MICHAEL O. LITTLE
a. REPORT	b. ABSTRACT	c. THIS PAGE			19b. TELEPHONE NUMBER (*Include area code*)
U	U	U	UU	23	N/A

TABLE OF CONTENTS

LIST OF FIGURES

Figure **Page**

1.0 SUMMARY

We have developed a theory for recovering clusters of important information from sparse signal representations in distributed sensing applications. This theory is expected to be useful in streamlining networking and decoding operations over bandwidth constrained wireless networks. Using these results, we can for example transport and store a single combined measurement set, rather than multiple sets from all sensors. We show that via source separation and joint decoding, it is possible to recover essential information of the original signal from such combined measurements using progressive reconstruction which, in each step, relaxes variables related to only a single sensor. This progressive procedure results in a substantial reduction in the number of variables that are decoded at each step, and consequently, a much reduced decoding time. We show that the reconstructed signal can still have sufficient accuracy for some target detection tasks. We demonstrate these results with image recovery for simple target detection examples.

We have documented this work and related results in several papers, including:

- “Separation-based Joint Decoding in Compressive Sensing,” 20th IEEE International Conference on Computer Communication and Networks (ICCCN 2011), August 2011.
- “Measurement Combining and Progressive Reconstruction in Compressive Sensing, Military Communications Conference (MILCOM 2011), November 2011.
- “Partitioned Compressive Sensing with Neighbor-Weighted Decoding,” Military Communications Conference (MILCOM 2011), November 2011.

2.0 INTRODUCTION

A compressive sensing encoding system can be described as follows [1] :

$$y = \Phi x \tag{1}$$

where x is an N-dimensional signal being sampled, Φ is an $M \times N$ measurement matrix containing random entries, and y is a vector of M measurements. We see that compressive measurements of y are random linear combinations of components of x. Suppose that x is K-sparse in the sense that it can be represented as a linear combination of K basis vectors in some basis, that is, $x = \psi s$ where ψ is an orthonormal transform, and s is a vector with no more than K nonzero coefficients. Then, if $M > cK \log (N/K)$ for some small constant c, it is possible to decode s with high probability, and recover $x = \psi s$. Note that the number of measurements (M) is commensurate with the sparsity K of a signal. If K is very small in some properly chosen domain, then M can be much smaller than N. This implies high compression rates. A rich volume of literature examines the topic of compressive sensing, including the landmark work of Candes and Tao [2] and Donoho [1] .

In this work, we address the use of compressive sensing in distributed sensing applications. In such a system, each sensor may gather compressive measurements for a specific region of a partitioned domain. That is, these distributed sensors perform partitioned encoding–one set of compressive measurements per partition–rather than standard encoding, as depicted in the middle and left sections of Figure 1, respectively. However, the aggregate size of all the measurements is proportional to the number of sensors and can result in large costs of transporting and storing these measurements, as well as a large decoding time for recovering signals from compressive measurements.

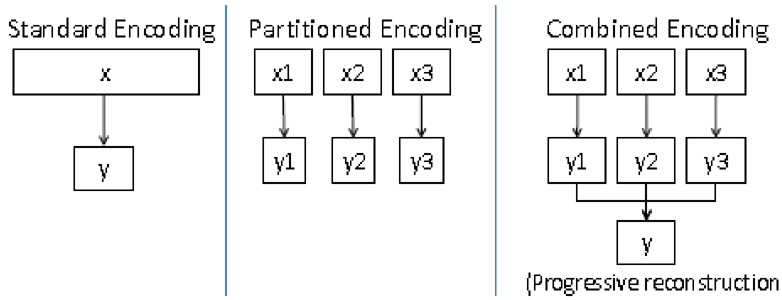

Figure 1: Methods of Encoding - (Left) Standard encoding, where the given signal x covers the entire domain. (Middle) Partitioned encoding, where partitions of x cover various regions of the domain. (Right) Combined encoding merges the measurements from all partitions to create a single measurement set

To reduce the number of measurements, we consider in this work a combining approach, as depicted by the Combined Encoding scheme of Figure 1. We show that we can combine multiple measurement sets from different sensors, by simply adding together the corresponding measurements to form a single measurement set. We can then transport and store a single measurement set, rather than measurement sets of multiple subproblems as depicted in Figure 2. In this example, there are 9 regions. For each region, we assume that there is one sensor encoding target information in that region, and sending compressed measurements to a hub to be combined. The hub combines the compressed measurements and transmits them to a remote decoder over a low-bandwidth link which benefits from measurement compression/combining.

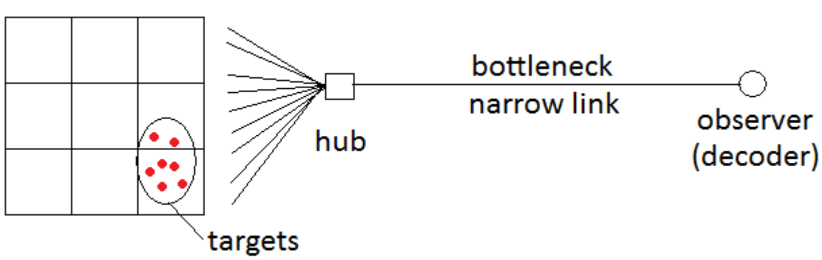

Figure 2: Partitioned scenario of tracking targets in an area divided into multiple regions.

In measurement combining, we will use a single measurement set to recover multiple sets of variables–one for each sensor in the distributed system. Since there are multiple sensors, this increases the problem size multiple times, and leads to another issue: high decoding cost. We note that the computational cost of the L_1-minimization step in decoding grows superlinearly with the problem size. If the signal has N components, then L_1-minimization by linear programming [2] can yield a decoding time of at least $O(N^3)$. This means that the decoding cost can be prohibitively expensive when N is large. Other decoding methods such as matching pursuit [3] , [4] are less expensive, but they only guarantee weaker error bounds in the recovered solution, and provide less reliable recovery. When these low-complexity methods fail, we still need to resort to L_1-minimization. It is therefore desirable to divide the problem into subproblems, each focusing on just one region at a time to reduce decoding time.

We show that we can reconstruct an approximate recovery for the original signal from the combined measurements by progressively decoding the subset of variables associated with just a single sensor at a time. This can significantly reduce decoding time in distributed sensor systems, as we will see later. Further, we demonstrate that the reconstructed signal resulting from progressive reconstruction can still have sufficient accuracy for target detection tasks.

Measurement combining and progressive reconstruction together can therefore reduce both the number of measurements and decoding time in distributed sensor systems, while yielding enough decoding accuracy for target detection. The approach is applicable to various application scenarios. For instance, as illustrated later, when an area is monitored by multiple sensors, each sensor may monitor its own nearby region for targets. Then we can use the method of this report to combine the measurements from such sensors, decode the combined measurements progressively, and finally detect targets from the reconstructed signal.

The measurement combining and progressive reconstruction ideas are based on some of our earlier theoretical work [5] . There we showed that via source separation and joint decoding, it is possible to separate out from combined measurements distinguished signal components in subproblems. The current work focuses on applications of these ideas in distributed sensor systems, and illustrates the combining approach with some simple image recovery examples for target detection.

3.0 METHODS, ASSUMPTIONS, AND PROCEDURES

We describe the measurement combining method in detail using an illustrative scenario. Suppose that we want to detect K targets from compressed measurements taken over regions, as depicted in Figure 2. Each sensor obtains a source signal vector x_i of length N from its own region, and then applies an $M \times N$ measurement matrix Φ_i with $M << N$, to obtain a compressed measurement vector $y_i = \Phi_i x_i$ of length M. The sensor then sends y_i to a hub node where the measurements may be combined.

3.1 Design Space

We describe a design space of methods that our system could use to reconstruct the signals of all regions. The first two are conventional methods that will serve as baselines for comparison purposes, while the last one is the focus of this work.

Conventional Decoding. Suppose measurement matrices Φ_i are all chosen independently at random. Furthermore, suppose the hub computes a sum y of the measurements y_i before sending them out; then, we can write the sum as follows:

$$y_1 + y_2 + \cdots + y_n = [\Phi_1 \quad \Phi_2 \quad \cdots \quad \Phi_\propto] \begin{bmatrix} x_1 \\ x_2 \\ \vdots \\ x_\propto \end{bmatrix} \tag{2}$$

We can let $y = \Sigma y_i$, use x to denote the column vector on the right-hand side of Eq. (2), and let Φ be an $M \times \alpha N$ matrix which is the column concatenation of $M \times N$ matrices $\Phi_1 \ldots \Phi_\alpha$. Then we have

$$y = \Phi x \tag{3}$$

This is in the conventional compressive sensing form. Thus, we can decode x using the standard formulation

$$y = \Phi \psi s \tag{4}$$

with a transform ψ operating on an αN-size vector of coefficients s. While this formulation requires a low number of measurements, $cK \log(\alpha N/K)$, it has a high decoding time. For example, if L_1-minimization is used, there can be an α^3-fold increase in the decoding time for each subproblem, $y_i = \Phi_i x_i$, due to an α-fold increase of the problem size from N to αN. The increase of variables is an issue we address in this work.

Partitioned Decoding. A simple way to avoid solving an αN-size problem instance is for the hub not to sum up the measurements y_i, but instead forward them directly to the decoder. Thus, the decoder ends up with α size-N compressive sensing problem.

$$y_i = \Phi_i x_i \tag{5}$$

This can be decoded much faster using the form

$$y_i = \Phi_i \psi_i s_i \tag{6}$$

where the transform ψ_i now operates on size-N vectors. However, this formulation requires that the hub transmit a larger amount of data; specifically, the total number of measurements to be transmitted is $\alpha cK \log(N/K)$. (Note that if fewer than $cK \log(N/K)$ measurements are used per region, then the signal may not be decodable if all K targets happen to be in the same region.)

Progressive Reconstruction. This is the decoding method we propose in this work. As in conventional decoding, the hub will again compute the sum of measurements $y = \Sigma y_i$, as in Eq. (2). How can we decode it faster than solving a size-αN problem? To answer this, let us consider the following decoding setup:

$$y = [\Phi_1 \quad \Phi_2 \quad \cdots \quad \Phi_\propto] \begin{bmatrix} \psi_1 & 0 & \cdots & 0 \\ 0 & \psi_2 & \cdots & 0 \\ 0 & 0 & \ddots & 0 \\ 0 & 0 & 0 & \psi_\propto \end{bmatrix} \begin{bmatrix} s_1 \\ s_2 \\ \vdots \\ s_\propto \end{bmatrix} \tag{7}$$

This is, again, a size-αN problem with a high decoding cost; however, note that each ψ_i here is α-times smaller than the one used in Eq. (3). We decode s_i separately using the constraint

$$y_i = \Phi_i \psi_i s_i \tag{8}$$

Note that this constraint is approximate since it ignores nonzero contributions of $s_{j\neq i}$ to y [5] ; nevertheless, in certain applications such as our target detection scenarios with sparse clusters, it is expected that these ignored contributions will be relatively small. Furthermore, as soon as larger s_i's are decoded for a region, we can use them to represent the entire region. This reduces the total number of variables in the subsequent decoding of other regions. The quality of each decoding step improves as more regions are decoded and represented with a reduced number of variables. Thus we call the method progressive reconstruction.

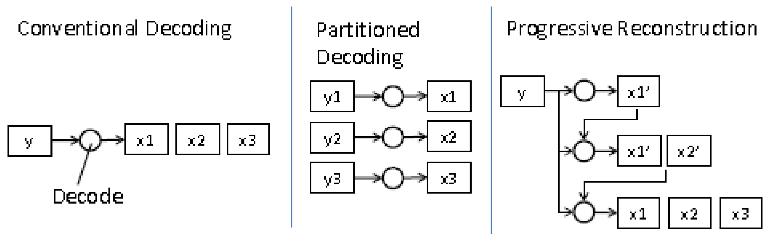

Figure 3: Methods of decoding compressed data - (Left) Reconstructing all regions of the combined measurement at once. (Middle) Decoding each partition separately. (Right) Progressive reconstruction focusing on a single region at a time while considering contributions from a few of the most significant variables in previously decoded regions.

Figure 3 depicts the three approaches. We note that conventional decoding involves a high decoding cost, due to a large number of variables. Partitioned decoding has a reduced cost, but at the expense of a poor compression rate. So, we propose the third solution—progressive reconstruction. Our method has a relatively high compression rate and low decoding cost, while achieving a decoding accuracy sufficient for target detection. In the rest of the report we describe progressive reconstruction in detail, and evaluate its performance.

3.2 Progressive Reconstruction on Combined Measurements

For simplicity, we describe our progressive reconstruction method for a partitioned problem involving only $\alpha = 3$ subproblems. Note that the description generalizes to any number of subproblems. Suppose that the signals for the three regions are: x_1, x_2, and x_3. As described earlier, we encode them with the measurement matrices Φ_i to yield three measurement sets, $y_i = \Phi_i x_i$, for $i = 1, 2$ and 3. We combine the y_i's to form y:

$$y = y_1 + y_2 + y_3 \tag{9}$$

In our scenario we assume that the targets are clustered, meaning that the large variables are concentrated in only a few of the regions.

The first step of the reconstruction method is to identify the region that contains the largest number of targets. This region can be chosen according to prior knowledge; for example, we might be able to predict a target's location based on past observations. In other applications where no such side information is available, the best region can still be determined by observing preliminary decoding results. A generic method works by decoding the combined measurements as follows:

$$s_i^* = \arg min\, |s_i|_{l_1} \quad \text{subject to } \; y_i = \Phi_i \psi_i s_i \tag{10}$$

for each region i. In the best region, by definition, the energy will be more concentrated on a few coefficients, allowing us to find a sparser solution. Given sufficient measurements, we will have enough constraints so that a sparse solution for other regions is unlikely. Therefore, we identify the region with the smallest $|s_i^*|_{l_1}$ to be the best region.

Finding this region is important because y has contributions from all three x, and each s_i^* may contain interference from other regions. Of these, the signal x_i corresponding to a region that has the most targets will experience the least interference from other sources, and will have the best reconstruction result.

We can now begin the progressive decoding stage. Suppose we find that the targets are concentrated in region 1 from the result of the first step. We then identify the basis associated with the largest variables in the reconstructed s_1^*, denoted as $\hat{s}_1$, and use it to represent x_1 in the subsequent steps. Specifically, we form $\hat{\psi}_1$ by keeping only the columns associated with the larger coefficients, as shown in Figure 4.

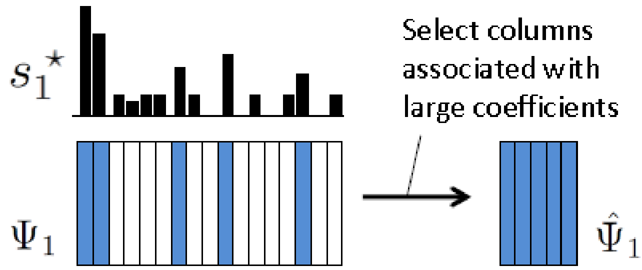

Figure 4: Forming a reduced basis for region 1

We continue the progressive decoding by reconstructing s2 in the second round as follows:

$$s_i^* = \arg min \left\| \begin{bmatrix} \hat{s}_1 \\ s_2 \end{bmatrix} \right\|_{l_1} \quad \text{subject to} \quad y = [\Phi_1 \hat{\psi}_1 \quad \Phi_2 \psi_2] \begin{bmatrix} \hat{s}_1 \\ s_2 \end{bmatrix} \tag{11}$$

Since $\hat{s}_1$ includes all the large variables in s_1, there is little or no interference from x_1 when we reconstruct s_2. Similar to the first round, we find the most important basis for s_2, and proceed to decode s_3 in the same fashion.

In each round, we slightly increase the number of variables to decode by including the basis vectors that are deemed to be important from all previous rounds. If no round misses important basis columns for the regions in question, then at the end of this process we will be solving an interference-free signal.

Compared to the progressive reconstruction process, conventional decoding which decodes signals from all regions in one shot exploits the sparsity better, as it represents the signal with a larger selection of basis vectors (a larger dictionary). However, in a real-world application like target tracking, we might have knowledge about the locations of the targets from a previous state. For example, if we know that most targets are in region *i* at a previous time, it is reasonable to assume that region *i* and its surrounding regions are more likely to contain the targets. In the progressive reconstruction method, we could take advantage of such knowledge by having early rounds focus on these regions.

4.0 RESULTS AND DISCUSSION

4.1 A Test Scenario

We consider a simple example in distributed sensor systems to demonstrate the idea of measurement combining and progressive reconstruction.

Suppose we want to track K targets across 9 regions, as depicted in Figure 2. The sensors compress the sensed images x_i by computing $y_i = \Phi_i \, x_i$. A hub gathers the measurements y_i and performs a simple summing operation to get $y = y_1 + y_2 + \ldots + y_9$, and then transmits y back to us.

In this example, we detect one or more copies of the same bee-like target as shown in Figure 5. To identify targets, we build a dictionary where each basis function represents the target at a distinct position in the region, as shown in Figure 5. More precisely, in the dictionary matrix, each column is a vector representation of the input image with the target at a distinct location. An image with k targets can be represented by k coefficients in this dictionary matrix. We can view the input signal x as the sum of those images, each containing a single instance of the target at a distinct location. Recovering x means finding these image components, or a vector s such that

$$x = Ds \tag{12}$$

where D is the dictionary matrix defined above. When there are just a few targets present in the signal, s is sparse and this problem will be solvable by compressive sensing decoding using the dictionary as the basis ψ_i described earlier. This is similar to finding parameterized shapes in images as described in [6] . For simplicity, we assume the background can be subtracted from the measurements.

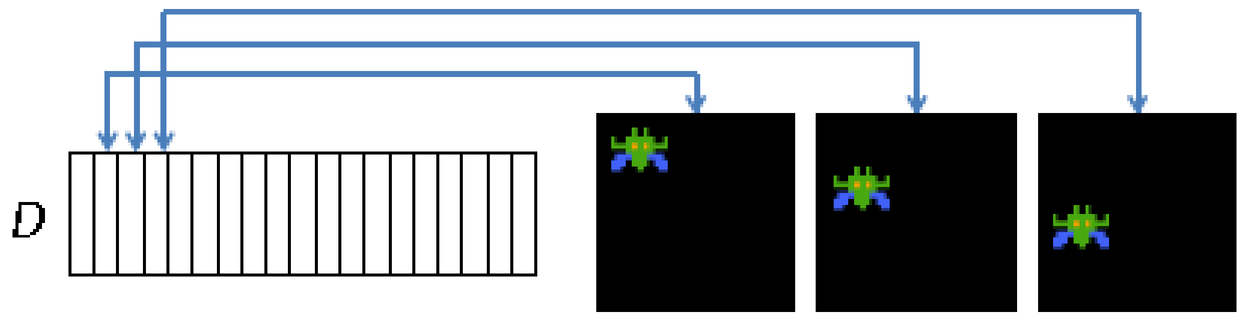

Figure 5: Dictionary matrix representation.

As described earlier, we perform a reconstruction for each region under some noise model and determine where the targets are most concentrated. Then, by progressively decoding one of the remaining regions at a time, we gradually improve the reconstruction until all regions are decoded. The whole process takes at most 9 + 8 + 7:: + 1 decoding steps, where there are roughly N variables to solve each step. Suppose that `1-minization is used in decoding. Then, the total cost is $45CN^3$, for some scaling constant C. In contrast, if we were to decode the entire image ($9N$) at once, the cost would be $C(9N)^3 = 729CN^3$.

4.2 Experiments and Performance Results

We examine the performance of progressive reconstruction in three cases:
1) Four targets spread across two regions;
2) Five targets spread across three regions;
3) Five targets spread across three regions with Gaussian noise (SNR $\approx$ 13).

The size of each region is 30 × 30 pixels. The three progressive cases each have 8100 pixels in total, whereas a fourth case (for comparison with case 3) uses standard decoding and has 5400 pixels. 60 measurements are used for the first two cases (compression ratio < 1%), 405 measurements are used for the third case (compression ratio = 5%), and 1200 measurements are used for the fourth case (standard decoding). We use L_1-magic as the decoder for the first three cases, and CoSaMP [4] for the last case.

Case 1
To begin progressive reconstruction, we must first find a region to start with. In the first round, shown in Figure 6, we decode the nine regions independently. Figure 6 (a) shows the decoded

results for each region decoded independently. The center region has the least interference from the other regions (displays pronounced targets), resulting in a better reconstruction. It is observed that the L_1-norm of this region is also smaller. Thus, we start the second round based on the reconstruction of the center region. In the second round, we reconstruct the signal in two regions together using a reduced basis for the center region as described earlier. We continue to recover other regions until all signals are decoded. Figure 7 shows three snapshots of the decoding process. The first decoding round revealed that the targets are most concentrated in the middle region. Then the second round of decoding uses the information from the first round to improve the reconstruction. Notice that since there are no targets left in other regions, the reconstruction in the second round is an exact reconstruction. This is due to the lack of interference from other regions.

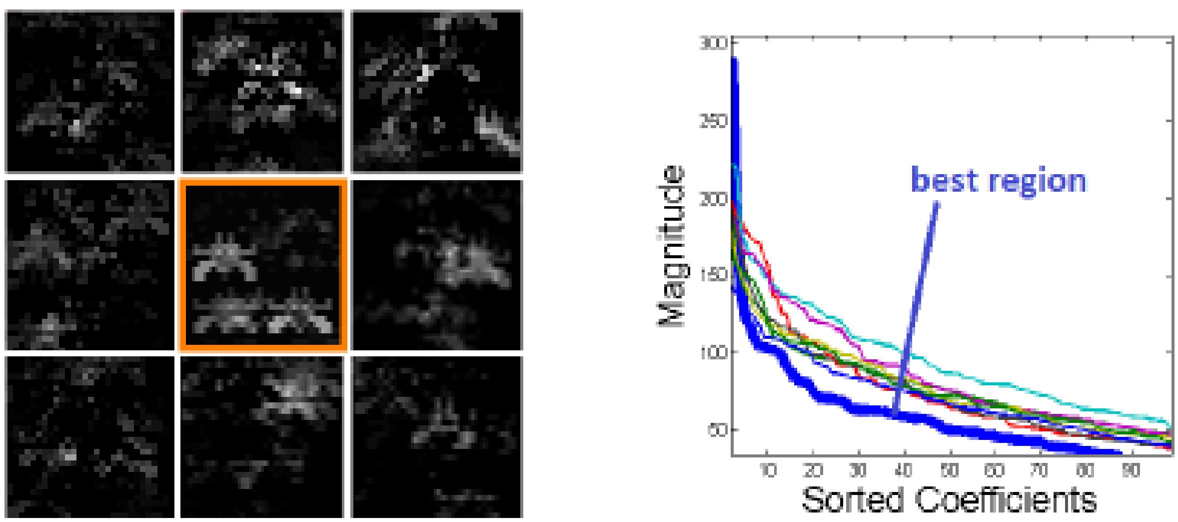

(a) Reconstructed image

(b) Magnitude of reconstructed coefficients for 9 regions

Figure 6: First decoding round for Case 1 - (a) The reconstructed images for each region. (b) The sorted magnitude of reconstructed coefficients for regions.

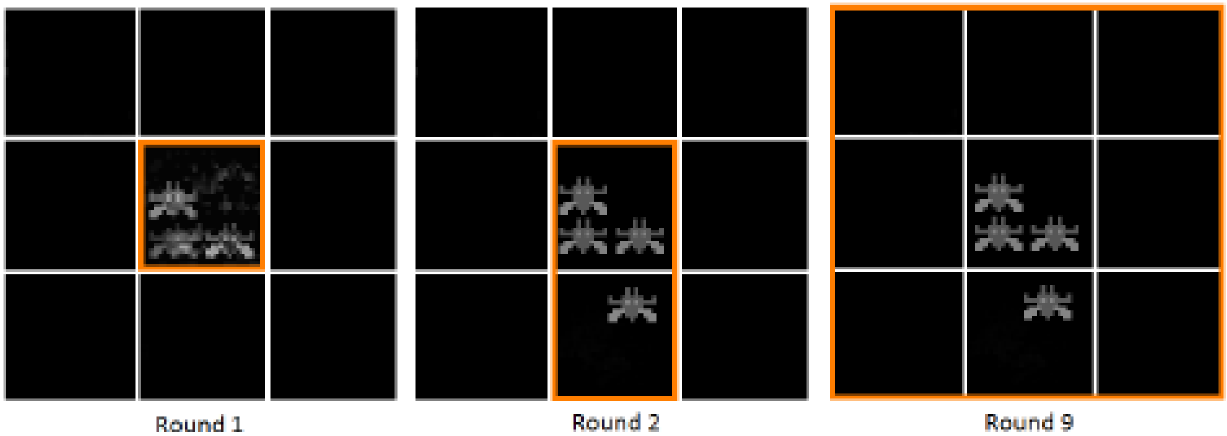

Figure 7: Progressive reconstruction for Case 1.

Case 2

Here we use the same decoding method as Case 1, but now there are three regions that contain targets. Figure 8 shows the first three rounds of the decoding process. As in Case 1, the first decoding round reveals that the targets are most concentrated in the center region. However, the reconstruction in round 1 is worse than the previous case, because now there is interference from the two additional targets located in two other regions. As in Case 1, the second round of decoding uses the information from the first round to improve the reconstruction. Since there is interference from a target that has not yet been detected, the reconstruction in the second round is still noisy. The reconstruction is improved again in the third round. Notice that the interference is progressively reduced until it falls to zero, and in this case, the reconstruction becomes exact after the third round.

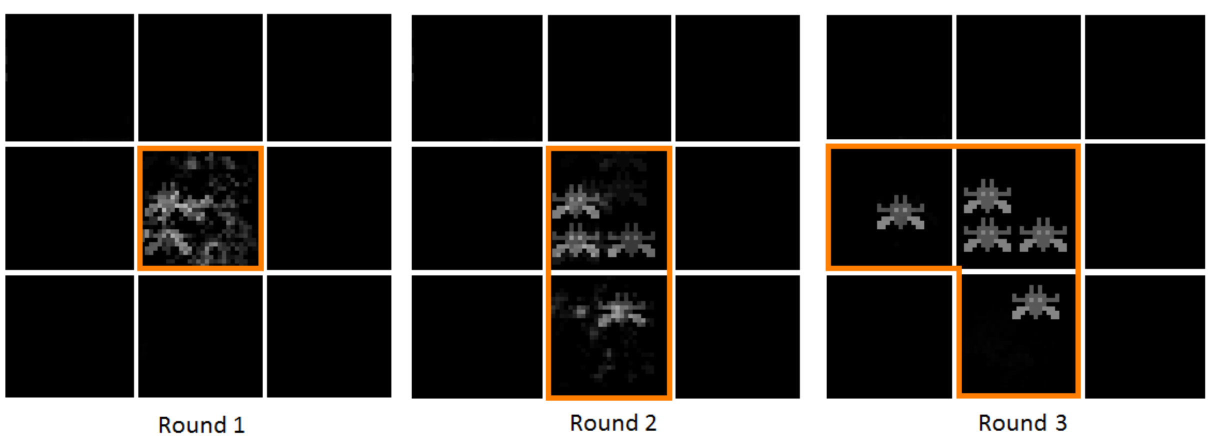

Figure 8: Progressive reconstruction for Case 2.

Case 3

In this case the signal is corrupted by noise. Figure 9 shows the results from each of the 9 rounds of the progressive reconstruction. Note, we focus on one region each round of decoding. The recovery in the first round is heavily distorted, but the three targets in the center region are still distinguishable. We only keep the bases associated with the largest 2% of the coefficients (refer to Figure 4). Note that distortion is most significant in the newest regions being solved; this is because we used reduced bases to represent the signals for regions that have already been processed, whereas the latest region added to the decoding is represented with a full dictionary. The results of progressive reconstruction and standard decoding (case 4) are compared in Figure 10. Both approaches locate the targets successfully, but the error from noise is distributed in a different way. While providing comparable results, the progressive reconstruction process is 10 times faster, assuming O(N^3) decoding time.

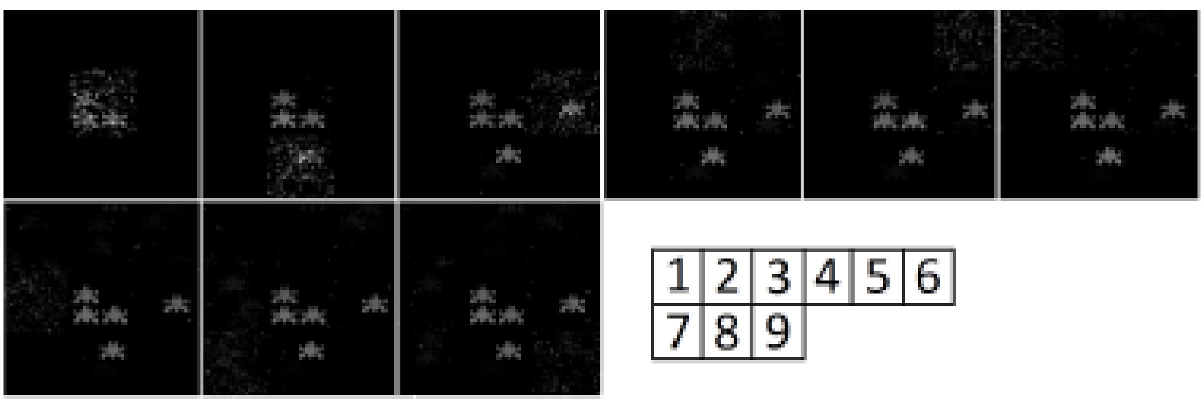

Figure 9: Progressive reconstruction for Case 3, a noisy image.

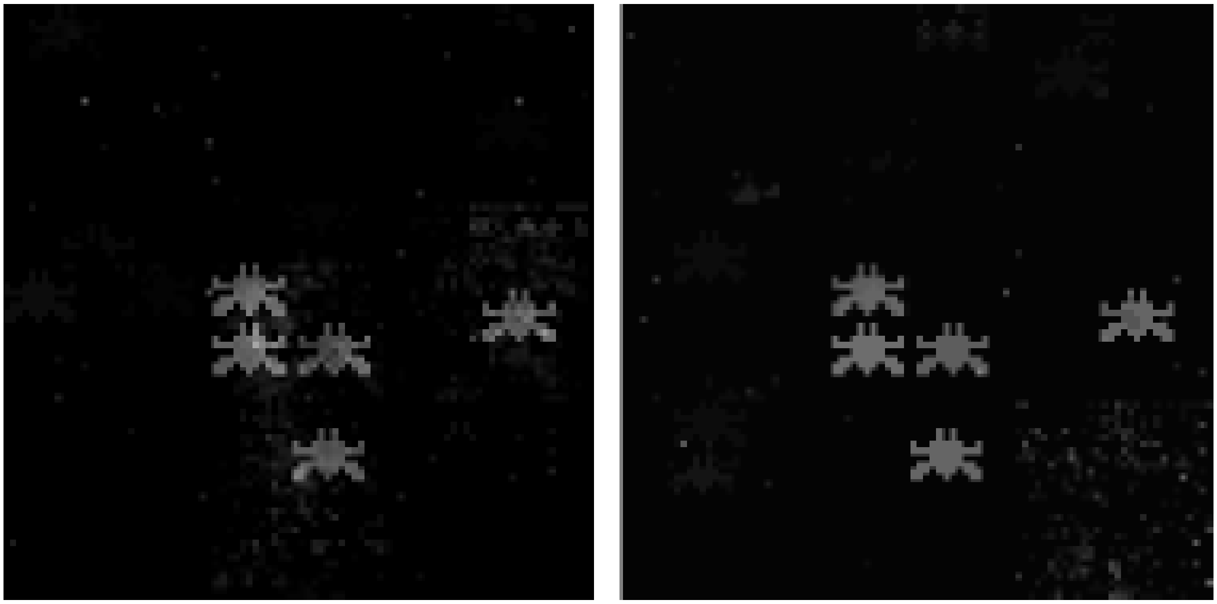

Figure 10: Comparison between progressive decoding (right) and decoding all regions (left) at once

4.3 Discussion

In the experiments, the reconstruction process can take advantage of additional side information. If we know the targets are clustered in a small set of regions, then our method can first identify the most populated region, and improve the reconstruction progressively at a low cost. Furthermore, if we know where these targets are most likely to be, this information can help us find a better reconstruction sequence and thus save more time. This can be a reasonable assumption for some surveillance systems where the targets are constantly monitored. When a number of measurements used is relatively small, there could be blocking artifacts resulting from the progressive reconstruction. In this case we can further apply methods such as partitioned compressive sensing with neighbor-weighted decoding [7] to ensure that the recovered signals for neighboring partitions will be compatible. Compressive sensing for distributed sensor systems is also studied in the context of collaborative decoding at a central node, such as a UAV, where the decoded results for one sensor can enhance the decoding for other sensors [8] . There are other approaches that also directly benefit from side information. Reweighted approaches [9] , [10] can use side information to guide the decoding process to a more desirable solution. However, these methods only improve the reconstruction quality but not the decoding time.

5.0 CONCLUSIONS

In this work, we take a perspective that an approximate signal reconstructed from compressive measurements just needs to be accurate enough for the target detection problem at hand. To this end, we presented a method of combining measurements in compressive sensing in order to lower the number of measurements, or equivalently to improve the compression ratio. In addition, we presented a progressive reconstruction method that can decode multiple signals from combined measurements with reduced transmission and decoding costs. We validate these results with target detection test cases. Our approach can reduce both measurement and decoding costs, while producing decoding results with sufficient accuracy for target detection. Results of this work are applicable to target-detection applications based on compressive sensing for which problem partitioning is desired and distributed sensing systems where compressive measurements from multiple sensors can be integrated.

6.0 REFERENCES

[1] D. L. Donoho, "Compressed sensing," IEEE Transactions on Information Theory, vol. 52, pp. 1289–1306, 2006.

[2] E. J. Candes and T. Tao, "Decoding by linear programming," IEEE Transactions on Information Theory, vol. 51, no. 12, pp. 4203–4215, Dec. 2005.

[3] D. Needell and R. Vershynin, "Signal recovery from incomplete and inaccurate measurements via regularized orthogonal matching pursuit," IEEE Journal of Selected Topics in Signal Processing, vol. 4, no. 2, pp. 310–316, Apr. 2010.

[4] D. Needell and J. A. Tropp, "CoSaMP: Iterative signal recovery from incomplete and inaccurate samples," Applied and Computational Harmonic Analysis, vol. 26, no. 3, pp. 301–321, May 2009.

[5] H.-C. Chen and H. T. Kung, "Separation-based joint decoding in compressive sensing," in 20th IEEE International Conference on Computer Communication and Networks (ICCCN 2011), 2011.

[6] A. C. Gurbuz, J. H. McClellan, J. K. Romberg, and W. R. Scott, "Compressive sensing of parameterized shapes in images," IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP), pp. 1949–1952, 2008.

[7] H. T. Kung and S. Tarsa, "Partitioned Compressive Sensing with Neighbor-Weighted Decoding," in MILCOM 2011.

[8] H.-C. Chen, H. Kung, D. Vlah, D. Hague, M. Muccio, and B. Poland, "Collaborative Compressive Spectrum Sensing in a UAV Environment," in MILCOM 2011.

[9] R. Chartrand and W. Tin, "Iteratively reweighted algorithms for compressive sensing," IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP), pp. 3869–3872, 2008.

[10] E. Candes, M. Wakin, and S. Boyd, "Enhancing sparsity by reweighted `1 minimization," Journal of Fourier Analysis and Applications, vol. 14, no. 5, pp. 877–905, Dec. 2008.

7.0 LIST OF SYMBOLS, ABBREVIATIONS, AND ACRONYMS

CoSaMP	Compressive Sampling Matching Pursuit
SNR	Signal-to-Noise Ratio
UAV	Unmanned Aerial Vehicle